A Tag-Mediated Game Designed to Study Cooperation in Human Populations

Garrison W. Greenwood
Department of Electrical & Computer Engineering
Portland State University
Portland, OR 97207–0751 USA
Email: greenwd@ece.pdx.edu

***Abstract*—How and why cooperation develops in human populations is not known. The iterated prisoner's dilemma game provides a natural framework for studying cooperation growth in human populations. However, recent experiments with human subjects has exposed a number of serious flaws in virtually all of the game-theoretical models that have appeared in the literature. Indeed, some experiments suggest network reciprocity—thought to be essential for cooperation in human populations—may actually play no role whatsoever. In this paper we briefly review some human experiments that were conducted in the last three years. We then present preliminary results of a new tag-mediated model designed for studying cooperation in human populations. The model exhibits many characteristics found in the human experiments including assortment, which many researchers now believe is necessary for maintaining cooperation.**

I. Introduction

Cooperation is pervasive throughout nature, although how and why it developed remains an open question. It is particularly difficult to figure out how cooperation evolves in humans because complex factors come into play that are not as strong (or even absent) in other species. Cooperation among social mammals is primarily limited to relatives whereas in humans cooperation exists at much larger scales and frequently among non-kin. Cultural information, which is information acquired by imitation or formal learning, influences human cooperation levels and it has been suggested cultural adaption may provide some explanations [1].

Unfortunately cultural adaption is an abstract concept that is difficult to quantify and even more difficult to incorporate into mathematical models. Various other explanations for cooperation have been offered and fortunately they can be tested by analyzing field data or modeling. For instance, the well-cited paper by Nowak [2] identifies 5 rules that promote cooperation: *kin selection, direct reciprocity, indirect reciprocity, network reciprocity and group selection*. With respect to human cooperation studies network reciprocity has sparked interest. In particular, human cooperation in dynamic networks has been the focus of much of this recent work (e.g., [3], [4], [5]).

The main idea behind network reciprocity is simple. In well-mixed populations individuals interact with everyone else with equal probability. Humans typically do not exist in well-mixed populations but instead have a small circle of individuals they interact with regularly (friends, relatives, neighbors and so forth). Networks provide a natural mathematical framework for describing such relationships, which facilitates a game-theoretical study of cooperation. That is, nodes represent individuals and only nodes that share an edge—called "neighbors"—are allowed to interact. In the game-theoretical studies "interaction" usually means playing either a 2×2 pairwise *iterated prisoner's dilemma* (IPD) game with each neighbor or one round of a *public goods game* (PGG). Networks can either be *static*, where link locations don't change, or *dynamic* where links are periodically added, deleted or moved. Dynamic networks have attracted a lot of recent attention because models predict if network changes are done frequently enough then cooperation can persist.

Over the past few years numerous computer models have been proposed to help understand the genesis of human cooperation. Agents in these models play rounds of IPD and update their strategies periodically according to defined rules. In prisoner dilemma games defection tends to be the inevitable outcome. The goal of these models is to find conditions under which clusters of cooperators will persist.

Yet, despite all of this time and effort, arguably little progress has been made. (See [6] for a more detailed discussion.) One of the main reasons why is models are seldom validated. That is, the models are developed and predictions are made, but the models are not compared against field data to see if the underlying mathematics are correct. In fairness field data to study human cooperation has not been readily available. But very recently a number of comprehensive human experiments have been conducted to address this model validation problem—but in the reverse direction. Instead of observing humans and then creating a model, the newest approach is to formulate a human experiment to match the model. If humans act the same way the models do, then the models must be correct. Unfortunately—for the model designers, that is—humans aren't acting like the models predict they would!

This paper was written with two purposes in mind. First is to provide a survey of the recent game-theoretical human experiments on cooperation that have been conducted and to discuss why the existing models lack accuracy. Second is to present preliminary results from a tag-mediated model recently developed by the author. This model is more realistic in the sense its design was based on observations made in the human cooperation experiments. It uses tags rather than

978-1-4673-5311-3/13/$31.00 ©2013 IEEE

network links or proximity to identify playing partners making it more flexible and natural. Instead of adding or deleting links players can change partners by changing tags. Players also have the option of not playing against an opponent, which makes ostracizing studies possible. Moreover, strategy updates are not based on a neighbor's payoff difference but rather on their actions, which is what was observed in a recent human experiment with a large number of participants [7].

The paper is organized as follows. In Section II a survey of recent research involving dynamic networks and human cooperation experiments is described. My tag-mediated model is described in Sections III and IV. Section V discusses the efficacy of prior computer models of human cooperation and puts the tag-mediated model described in this paper into context. Finally in Section VI future research directions are discussed.

II. Recent Research

Over the last 3 years several experiments have been conducted using actual human subjects playing either an IPD or a PGG game. These experiments were constructed like previously published computer models to see just how accurate those models are. If the human experiments produced the similar results, then the computer models would be deemed validated. Those human experiments are briefly described in this section.

The conclusions drawn in these papers fall into two categories. The first set of papers conclude networks contribute little or nothing to cooperation levels. In other words, network reciprocity has nothing to do with promoting cooperation in human populations. The second set of papers conclude static networks do not contribute to cooperation levels but dynamic networks do make a strong contribution.

Grujic et al. [8] constructed a spatial model where 169 humans were placed on a (virtual) 13×13 lattice. During each round agents interacted with their 8 nearest neighbors (Moore neighborhood) in a 2×2 IPD game. Three types of population structures were investigated. In experiment #1 a static lattice was used so neighbors remained fixed throughout the entire game. In experiment #2 agents also were placed in a lattice but in this case the neighbors where shuffled each round. Finally, in experiment #3 the neighbors were fixed in the lattice but this time agents were told what actions their neighbors chose in the previous round and what payoff they received. Agents chose the same action (C or D) for all of their neighbors. In all three experiments the cooperation levels leveled off at around 20-25%. Since this level was reached regardless of whether the neighbors were fixed or changed each round, the authors concluded the network contributed nothing to the cooperation level in the population.

Two other interesting conclusions resulted from the Grujic et al. experiments. First, they looked at the "imitate-the-best" update rule. This update rule appears frequently in computer models. It assumes a focal agent adopts the strategy of his best fit neighbor with a probability proportional to the difference in accumulated payoffs. They found humans do *not* tend to follow the imitate-the-best update rule. They also found that while as high as 45% of the participants play C or D (almost) all of the time, the remaining participants tend to choose their next play based on their last choice and the last choice of their opponent. That is, accumulated payoffs were not part of the decision; only the choices made during the last round have any bearing.

Gracia-Larzaro et al. [7] conducted PD games with 1,229 human participants on lattices and scale-free networks. The authors found no significant difference in the cooperation levels between either of these networks and a well-mixed population. The authors therefore concluded network reciprocity has little to do with cooperation levels. They also found agents were more likely to cooperate if the majority of their neighbors cooperated in the previous round. These results support the Grujic et al. experimental results.

Traulsen et al. [9] analyzed update strategies in spatial networks. Four hundred students were arranged in a virtual lattice with Von Neumann neighborhoods. The first attempted to resolve whether "imitate-the-best" updates are used by human. Initially they found a high percentage (¿60%) of the participants did use imitation, but this decreased by around 4% per round with no stationary state encountered. One interesting outcome was the probability of cooperating versus number of cooperating neighbors. They found the highest probability if all neighbors cooperated—even though in such cases defection would produce a higher payoff. The authors concluded humans may imitate what is common instead of what pays the most. They also found that fixing the neighbors or randomly shuffling the neighbors had no real effect on cooperation levels.

Rand et al. [10] studied dynamic networks, which give agents another method of responding to their opponents last action. They used 765 participants in a (virtual) random network. Three forms of link updates were studied: random link updating, fixed links (no updating) and "strategic" link updating. In the latter case at the end of each round a small fraction of the agents (10% or 30%) are given an opportunity to remove existing links or add new links. Agents were informed of the action of the opponent in the previous round prior to adding to deleting a link. Each agent chose a single action (C or D) against all of his neighbors. In the random link and fixed link experiments the cooperation levels decreased asymptotically over time. In they dynamic network case the cooperation level also decreased when only 10% of the links were updated, but cooperation levels were robust and stable when 30% of the links were updated. They concluded cooperation can be maintained if network rewiring occurs at a high enough rate.

Wang et al. [11] used 108 human subjects in their PD game experiments. Participants were allowed to establish or break relationships at variable rates. They found updating relationships had a positive effect on cooperation. However, their results (and conclusions) are difficult to put in context with other related research work. The reason why is the Wang et al. experiments were constructed quite differently from the way others had done. For example, one observation they made was dynamic link updates produce high cooperation levels, but those levels decreased significantly near the end of the game. This behavior, they claim, deviates from the predicted Nash equilibrium of all defectors throughout the game. Could the fact the players knew the games were only 12 rounds have been a factor? This author is unaware of any other human experiment reported in the literature (going back almost 10 years) where participants were told how long the game would last. Experimenters specifically don't disclose that information to avoid unintentionally biasing the game play. Also the way Wang et al. recruited participants allowed some individuals to play several games while others played only one game. Asymmetric learning is therefore introduced into the Wang et al. experiments.

Fehl et al. [3] used 200 human subjects partitioned into two equal groups. The first group had fixed neighbors while the second group could change neighbors. In both groups each individual had 3 neighbors. Payoffs were awarded after each round of a PD game. Individuals knew the payoff their opponent received the previous round, but did not know the opponent's total payoff[1]. In the dynamic network at the end of each PD round partners were asked if they wanted to remain neighbors. Both had to agree or the network link was broken. In such a case both partners would receive new links taken from the set of individuals seeking new partners. This maintained the 3 neighbor requirement. They found dynamic networks yielded higher cooperation levels. They also found that individuals were much more likely to break a link with a defector than with a cooperator—i.e., $C-C$ links lasted longer than $C-D$ links. The authors believe this result is due, in part, to the limited neighbor size because maintaining a link with a defector limited future opportunities. In other words, breaking a $C-D$ link creates an opportunity to hopefully form a $C-C$ link in the future. They also found the more cooperative an agent is, the higher the group score is, which gives rise to a property, called *assortment*. This property is believed to be a crucial component for cooperation in populations [12].

III. A Tag-Mediated Model

Let N be the number of agents in a population. The agent genotype consists of $2N$ real numbers on the unit interval and a "tag". There are N $p_j(k)$ terms representing the probability agent j will cooperate with agent k. There are also N $q_j(k)$ terms representing the probability agent j will interact with agent k. The tag set for the population is $Tags = \{black \quad red \quad blue \quad yellow \quad green\}$. Each agent is assigned a single tag and he can only interact with other agents who have the same tag. That is, $q_j(k) = 0$ if agent j and agent k have different tags. Notice agents with the same tag do not necessarily interact during a given round because $q_j(k) \leq 1$. Of course, $p_j(k) = 0$ if $q_j(k) = 0$. Also, $p_j(j) = 0$ to preclude self-play.

[1]Of course in static networks these totals would be known, but not in dynamic networks since opponents are constantly changing.

Game play during each round involves the following steps. First, each agent compares his tag with the other $N-1$ agents. Only agents with matching tags can interact. Assume the tags of agent i and agent j match and let i be the focal agent. Then agent i agrees to play against agent j with probability $q_i(j)$ and agent j agrees with probability $q_j(i)$. Both must agree or no interaction takes place—even though the tags match. If they both agree, then agent i cooperates with probability $p_i(j)$ and agent j cooperates with probability $p_j(i)$.

Let C and D be represented by unit vectors

$$s = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

respectively. Then agent i's play is denoted by the unit vector s_i. When agent i plays agent j then the payoff to the focal agent i is $s_i^{\mathrm{T}} Q s_j$ where Q is the payoff matrix

$$Q = \begin{array}{c} \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cc} C & D \end{array} \\ \begin{pmatrix} R & S \\ T & P \end{pmatrix} \end{array} \tag{1}$$

with $T > R > P > S$ and $2R > (T + S)$. These are typical IPD relative values. (In this work $T = 0.4, R = 0.3, P = 0.1$ and $S = 0.0$ were used.) Only the focal agent i collects the payoff; agent j must wait until it's his turn to be the focal agent. Agents accumulate payoffs from every other agent with a matching tag and where mutual agreement to interact was reached. The round is completed when every agent has been a focal agent.

This model allows groups to ostracize exploitive agents but reward cooperative agents with even higher probabilities of cooperation in the future. Each agent j has a probability $p_j(k)$ of cooperating with agent k. This probability will either slightly increase or slightly decrease depending on what agent k did in the previous round when interacting with agent j. $p_j(k)$ is adapted as follows:

$$p_j(k) = \begin{cases} p_j(k) + \Delta p & \text{if agent } k \text{ cooperated} \\ p_j(k) - \Delta p & \text{otherwise} \end{cases} \tag{2}$$

where Δp is a random number on the interval $[0, 0.2)$. In addition, the interaction probability $q_j(k)$ is also adjusted according to what play agent k made in the previous round. This probability is increased (decreased) by 0.25 if agent k cooperated (defected) against agent j. (This probability is bounded on the interval $[0, 0.9]$.)

Adjusting the interaction probability allows a cluster of cooperative agents to ostracize a defective agent. Let M be the number of agents with the same tag and define $\xi = \lceil M/2 \rceil$. Suppose agent k is in this tag group. At the end of every round the number of agents in this tag group with $q(k) < 0.5$ are counted. If at least ξ of these agents do, then agent k is forced to leave the group. Agent k joints another group by changing tags. Agents in the old group set $q(k) = 0$ and $p(k) = 0$ while agents in the new group set $q(k) = 0.8$ and $p(k) = 0.45$.

IV. Experimental Results

The model described in the previous section was run with $N = 100$ agents. Initially these agents were evenly distributed among 5 tag groups: *black, red, blue, yellow* and *green*. Thus, the population started out with 20 agents in each tag group. The cooperation probabilities were initialized to $p_j(k) = 0.5$ if j and k have the same tags and 0 otherwise. Interaction probabilities were initialized to $q_j(k) = 1.0$ if j and k have the same tag and 0 otherwise. ($p_k(k) = q_k(k) = 0$ to preclude self-play.) These values insure the game begins with agents in the same tag group agreeing to play each other and with an equal probability of cooperating in the first round. Thereafter these probabilities would adapt based on the outcomes of pairwise encounters. The PD payoff matrix described in the previous section determined the payoffs, which were summed to get a player's accumulated payoff. Each game consisted of 70 rounds.

Figures 1–3 show the results of a typical run. As the game is played some agents are naturally ostracized and must join a new group. Since the group sizes are dynamic the average cooperation probability of the group will also be dynamic. Figure 1 shows the average probability of cooperation recorded per group size for each of the 5 tag groups. Although not strictly monotonic, the trend is quite clear: smaller group sizes tend to have a higher group average cooperation probability. Assortment is clearly present.

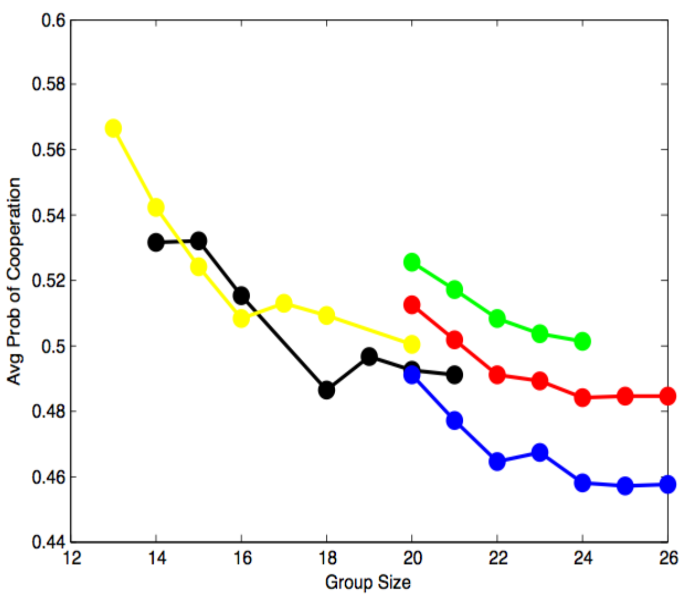

Fig. 1. Average probability of cooperation versus group size. Colors represent different group tags.

Agents that tend to choose D will ultimately be forced to switch tag groups. The number of groups agents belonged to is shown in Figure 2. Approximately 20% of the agents switch tag groups once, while one agent had to switch twice. The rest of the agents remained in their initial tag group. Of course this rate will change if ξ changes.

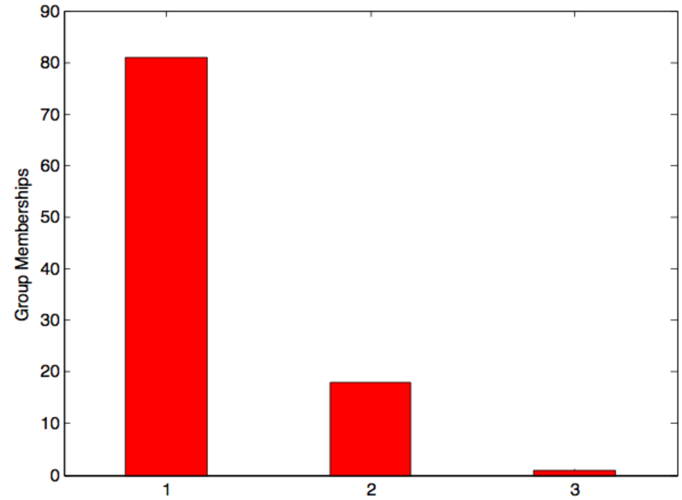

Fig. 2. Frequency of tag changes.

Figure 3 shows the distribution of agents after 70 rounds of a game. The red line shows the initial distribution size. The smaller groups have ostracized more agents. It is much easier to ostracize defectors if the majority of tag group agents have high cooperation probability—i.e., if positive assortment exists.

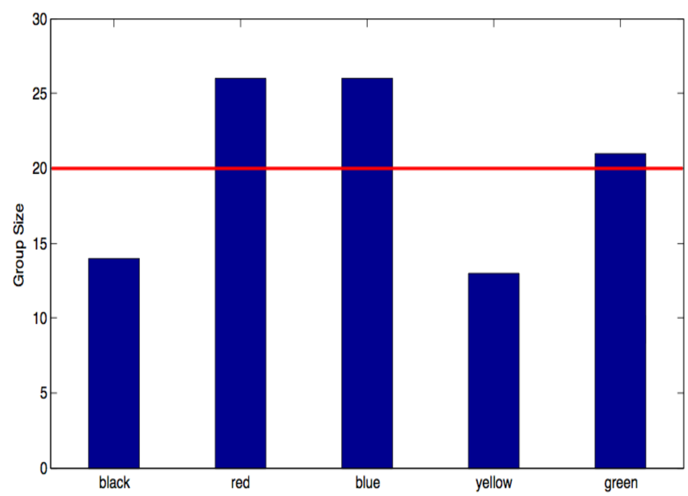

Fig. 3. Final distribution of agents among tag groups. The red line show the initial distribution.

Table I shows the group sizes and average group cooperation probability after the last round of a game. (This was the same game used to generate Figures 1–3.) The table values are slightly different from the data plotted in Figure 1. The table values are a snapshot taken at the end of the game whereas the figure values are averaged over all 70 rounds. There may have been multiple rounds during the 70 round game when the green tag group had 21 agents; the average over all of those particular rounds is plotted Figure 1 data. Conversely, Table I shows only one such instance—i.e., after the 70th round.

Table I also shows the interaction probabilities. Recall the

TABLE I
AVERAGE COOPERATION IN TAG GROUPS AFTER 70TH ROUND

Tag Group	Final Group Size	Avg Probability of Cooperation	Avg Probability of Interaction
Black	14	53.2%	57.8%
Red	26	48.0%	56.2%
Blue	26	45.7%	61.5%
Yellow	13	56.4%	62.2%
Green	21	51.8%	55.5%

interaction probabilities for at least ξ agents in a group must be less than 0.5 to ostracize an agent. Notice these probabilities are rather high, which explains why so few agents were ostracized.

Using the data from Table I we can estimate the average cooperation level in the population. Two agents in the same tag group will, on average, mutually cooperate with probability $q_{avg}(\cdot)p_{avg}(\cdot)$. Using the average probability values the intra-group cooperation probabilities tag group cooperation ranges from 26.8% to 35.1% and a population average of 29.9%. This matches well with the Gracia-Lazaro et al. [7] human experiments where they observed population cooperation levels settled to around 30% regardless of the network topology[2].

V. DISCUSSION

Over the past 20 years literally thousands of computer models have been created to help explain how cooperation evolves in nature. Arguably little progress has been made and we still don't have good answers. As discussed in [6], there are two reasons why this is so:

1) unvalidated models
2) unrealistic computer models

First consider the validation issue. The majority of papers discussing cooperation that include a computer model follow a familiar theme. The paper begins with an observation of a population in nature or the outcome of a human experiment, a finite population computer model will then be coded—and the remainder of the paper will methodically analyze the computer model. Claims will be made regarding the evolution of cooperation based on model predictions. But what the authors rarely do is verify their model outcomes match the observations or experimental outcomes before generating predictions. The second issue arises because models are invariably have a stochastic component. Many models random vary some parameter to simulate evolution. For example, a randomly chosen link in a network might be deleted or a new link randomly added. Either action changes agent neighborhoods. An agent may imitate a neighbor's strategy with some probability. The problem here is there is no rationale given for such actions. More importantly, recent studies have shown recently published computer models attempting to explain human cooperation are just plain wrong!

[2]Gracia-Lazaro et al. point out this 30% level is much lower than what computer models in the literature have predicted.

The flaws in previously published models of human cooperation were determined in a clever way. Normally one does (or at least should) do the following:

1) take measurements or observations of a physical system
2) create a mathematical (computer) model to describe the measurements or observations
3) compare the model output with the physical system measurements or observations

The model is "validated" if its outputs are in some sense ϵ-close to the physical system outputs. The model can then be used to make reliable predictions.

As stated previously, most researchers do not do this 3rd step but instead mathematically analyze their models for things like take-over times or invasion probabilities. Recently what researchers interested in human cooperation have done is validate the models in a reverse manner:

1) run simulations with an existing model to get model outputs
2) construct a human experiment patterned after the model parameters/dynamics
3) conduct the human experiment to get results
4) compare the experimental results with the model outputs

If the humans behave the way the models predict they would, then the model is validated. However, if the humans behave differently, then the models are flawed.

The problem with unrealistic models is more pervasive and may very well be the root of the problem on why computer models don't match human experimental results. For example, in the Grujic et al. [8] and Rand et al. [10] models all players were required to choose C or D against all of their neighbors. While that requirement does make the human experiment match the computer model, it is not clear how that contributes anything to understanding how cooperation occurs in human populations. Humans don't behave that way. What self-interested agent would continue to play C against an opponent who consistently plays D? Furthermore, defecting to punish a defecting neighbor also punishes a cooperating neighbor. The best outcome for a self-interested agent is to choose different strategies against different neighbors depending on how they previously played—i.e., direct reciprocity applied on an individual basis.

The model described here does not require a focal agent to play the same strategy (cooperate or defect) against all of his neighbors. Instead an agent j cooperates with agent k with some probability $p_j(k)$. A separate probability is kept for each agent in the same tag group, which means agents can play a different strategy for each neighbor. Unlike most other models a separate interaction probability $q_j(k)$ is kept for every neighbor. If $q_j(k) < 1$ then agent j doesn't necessarily play against agent k in the current round. This provides an opportunity for a cooperator to shun a defector—i.e., it permits punishment. By adapting this interaction probability it allows non-cooperating agents to be thrown out of the group. This feature leaves cooperators in the tag group thereby creating positive assortment.

It is interesting to note that in the extensive simulation that was run on the tag-mediated model in a few rare instances a tag group had a single agent. That says a cooperative agent would rather not participate than be continually exploited by a defector.

Over the past 10–15 years many models have used an imitation update rule for evolving strategies. Unconditional imitation means an agent will always switch to the strategy of an agent with a higher accumulated payoff. However, usually imitation is applied with some probability. A common method used mimics the Fermi formula from statistical physics. That method works as follows. Two individuals x and y are chosen. Let π_x and π_y be the fitness of x and y, respectively. Then with a payoff difference $\Delta\pi = \pi_x - \pi_y$, individual x replaces individual y with a probability

$$p = \frac{1}{1 + e^{-\beta\Delta\pi}}$$

where β is the selection intensity. If $\beta \to 0$ then switching strategies is random; as $\beta \to \infty$ the imitation is unconditional.

An interesting conclusion drawn from several human experiments is humans do not use an "imitate-the-best" update rule. As stated in [7], "... thus confirming that subjects did not use payoff differences as a guidance to update their actions". This conclusion is intuitively correct since no logical agent would not rely on payoff differences to decide whether or not to adopt a new strategy.

To better understand this concept consider the spatial game shown in Figure 4. Individual x may have a higher payoff than y, but that payoff was acquired from interacting with players individual y does not interact with. Put another way, the x and y neighborhoods are not the same. Just because a particular strategy works in $x's$ neighborhood does not mean it will work in $y's$ neighborhood. Hence, "imitate-the-best" is inherently flawed. This may very well explain why unconditional imitation was either not observed in human experiments or, when it was, it rapidly dissipated.

Kin selection, direct reciprocity, indirect reciprocity and network reciprocity have all been offered as reasons why cooperation grows in populations. A more general explanation for the evolution of cooperation is *assortment*, which describes how agents are distributed with respect to some characteristic. Cooperation flourishes under positive assortment where cooperators are more likely to interact with other cooperators than with defectors [12]. Such associations help offset any costs incurred by the cooperation. Extensive simulations show the tag-mediated model consistently develops higher average cooperation probabilities in the smaller tag groups. This could only happen if the majority of the tag group members cooperate with a high probability—a state that is much easier to reach in smaller groups.

The tag-mediated model promotes assortment by increasing the interaction probability and the cooperation probability when agents cooperated in the last round. Increasing the probabilities in this manner incorporates a Gracia-Lazaro et al. experiment conclusion that humans modify their strategies based almost entirely on the last action of their partner and not from some extensive history or reputation (indirect reciprocity).

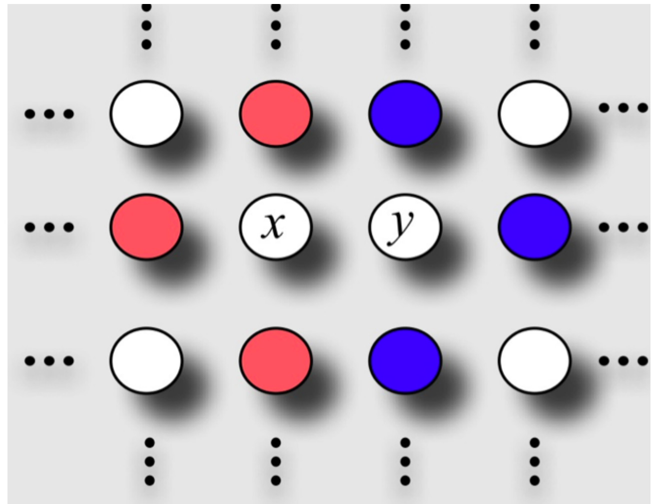

Fig. 4. A Von Neumann neighborhood in a spatial game. The red (blue) nodes are cooperators (defectors). Nodes x and y are in each other's neighborhood.

VI. Final Remarks

The tag-mediated model is general enough and flexible enough to make it the preferred model framework for studying cooperation in human populations. Some the advantages include

- *Neighborhood Membership*
  In some human experiments the neighbors were fixed while in others they changed every round. Agents can only interact with other agents with the same tag. Fixed neighborhoods are created by assigning the tags and never changing them. This is easily accomplished by setting the interaction probabilities $q_j(k) = 1$ if agents j and k have the same tag and never adapting those probabilities. This would prevent ostracizing agents thereby keeping the membership intact. Random shuffling of neighbors is accomplished by randomly reassigning tags each round.
- *Network Structure*
  The spatial networks used in some human experiments are of dubious value. Humans do not naturally associate in lattices but they do associate in either random networks or scale-free networks. The former can be created by randomly assigning the tags. The latter can be created by assigning agents more than one tag. Multiple tags allows agents to belong to more than one tag group. A scale-free networks is then created by making a few tag groups very large while making most of the other tag groups relatively small.
- *Dynamic Networks*
  Dynamic networks are created by adapting the interaction probabilities. The frequency of network changes can be controlled by adjusting ξ.

One area that we intend to investigate in the future is group competition. In a PGG environment free-riders will tend to decay cooperation within a group. However, when competition between groups exists studies suggest within group cooperation increases thereby resolving the social dilemma. Recently a human experiment was conducted to study this concept [13]. In that experiment groups were shuffled each round, which prevented seeing what roles punishment or guilt might play. Our tag-mediated model can ostracize non-cooperating agents. It will be interesting to see how that possibility affects within group cooperation levels.

REFERENCES

[1] R. Boyd and P. Richerson. Culture and the evolution of human cooperation. *Phil. Trans. R. Soc. B*, 364:3281–3288, 2009.

[2] M. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006.

[3] K. Fehl, D. VanderPost, and D. Semmann. Co-evolution of behavior and social network structure promotes human cooperation. *Ecol. Ltrs.*, 14:546–551, 2011.

[4] M. Zimmermann, V. Eguiluz, and M. Miguel. Coevolution of dynamical states and interactions in dynamic networks. *Phys. Rev. E*, 69:065102–1–065102–4, 2004.

[5] N. Hanaki, A. Peterhansl, P. Dodds, and D. Watts. Cooperation in evolving social networks. *Manage. Sci.* , 53(7):1036–1050, 2007.

[6] G. Greenwood and D. Ashlock. Evolutionary games and the study of cooperation: why has so little progress been made? In *Proc. 2012 IEEE Congress on Evol. Comput.*, pages 680–687, 2012.

[7] C. Gracia-Lzaro, A. Ferrer, G. Ruiz, A. Tarancn, J. Cuesta, A. Snchez, and Y. Moreno. Heterogeneous networks do not promote cooperation when humans play a prisoners dilemma. *Proc. Nat'l. Acad. Sci.*, 109(32):12922–12926, 2012.

[8] J. Grujic, C. Fosco, L. Araujo, J. Cuesta, and A. Sanchez. Social experiments in the mesoscale: humans playing a spatial prisoner's dilemma. *PLoS ONE*, 5(11):doi:10.1371/journal.pone.0013749, 2010.

[9] A. Traulsen, D. Semmann, R. Sommerfeld, H. Krambeck, and M. Milinski. Human strategy updating in evolutionary games. *Proc. Nat'l. Acad. Sci.*, 107(7):2962–2966, 2010.

[10] D. Rand, S. Arbesman, and N. Christakis. Dynamic social networks promote cooperation in experiments with humans. *Proc. Nat'l. Acad. Sci.*, 108:19193–19198, 2011.

[11] J. Wang, S. Suri, and D. Watts. Cooperation and assortativity with dynamic partner updating. *Proc. Nat'l. Acad. Sci.*, 109(36):14363–14368, 2012.

[12] J. Fletcher and M. Doebeli. A simple and general explanation for the evolution of altruism. *Proc. R. Soc. B*, 276:13–19, 2009.

[13] M. Puurtinen and T. Mappes. Between-group competition and human cooperation. *Proc. R. Soc. B*, 276(1655):355–360, 2009.